



Abuse of Customer Premise Equipment and Recommended Actions

Paul Vixie,[†] Chris King,* Jonathan Spring*

Netsa-contact@cert.org

CERT[®] Coordination Center, Software Engineering Institute

[†]Farsight Security, San Mateo, CA.

* CERT[®] CC, SEI, Pittsburgh, PA.

Publication CERTCC-2014-48

August 7, 2014

Executive Summary

Customer Premise Equipment (CPE) connects the customer's network to the service provider. This used to be as simple as a Bell Atlantic telephone, but in the modern age it includes many device types. Focusing on home Internet routers, we demonstrate some of the present dangers of the current CPE environment and possible solutions.

Namely, CPE can be used by adversaries to amplify and anonymize their denial of service (DoS) attacks, and the CPE itself can be compromised as part of an attack to redirect the customer's Internet traffic for illicit gains. The scope of this problem is large: Of the 22 million open domain name system (DNS) resolvers connected to the Internet as of May 2014, the majority are on connections indicative of home Internet users.

Misconfigured or outdated routers and CPE present essentially a public health hazard to the Internet. The poor digital hygiene of these devices (relatively few in the scope of the Internet) threatens the general enjoyment of the resource for any given Internet user, given the ease with which the misconfigured CPE can be abused to amplify attacks.

In order to counter this threat, we present three recommendations: (1) provide for continuous software upgrades of CPE, (2) implement source address validation (i.e., Best Current Practices document 38 and/or 84), and (3) encourage the community to incentivize manufacturers and providers to take responsibility for the results of poor configuration and design choices.

Copyright 2014 Carnegie Mellon University

This material is based upon work funded and supported by Department of Homeland Security under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by Carnegie Mellon University or its Software Engineering Institute.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution except as restricted below.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

* These restrictions do not apply to U.S. government entities.

Carnegie Mellon®, CERT® and CERT Coordination Center® are registered marks of Carnegie Mellon University.

DM-0001510

1. Introduction

The modern Internet ecosystem is complex. Like in any complex ecosystem, small details can have large unexpected consequences. Most of these unexpected consequences can be ignored. However, sometimes one error propagates to a tipping point, and it becomes a systemic concern. On the Internet, small errors that can be reliably targeted in many devices can lead to systemic security threats that could be abused to threaten public enjoyment of the resource. We believe that misconfigurations and vulnerabilities in customer premise equipment (CPE) such as home wireless routers has approached such a dangerous point. Reining in these errors will take concerted effort from device manufacturers, network operators, and end users. First this paper will provide some measurements of the problem and then provide some recommendations.

1.1. Customer Premise Equipment (CPE) background

Strictly speaking, CPE is any device located with the consumer of telecommunications services. As such, it provides an interface between the customer's local network and the telecommunication provider's network. CPE includes devices such as telephones and private branch exchanges (PBXs) to interface between an organization and the phone company, DSL routers to connect a home network to the public Internet, and set-top boxes to interface with the cable TV network. While all of these devices, as computers, may have flaws, the types of CPE that we are largely concerned with here are home routers, wireless routers, and modems that connect the Internet service provider (ISP) to the consumer's local network. These consumers include home users as well as small- and medium-sized businesses.

In the Internet ecosystem, the CPE devices may be managed and supplied by the ISP or by the consumer. There is no consensus that either ISP-managed or consumer-managed devices are better. Regardless, the diversity of the devices does provide a challenge. An additional challenge is that these devices do not have their own physical management interface, so many users do not know how to even begin to configure these complicated devices, let alone do so correctly. From a security perspective, this makes default configurations particularly important, but in the competitive marketplace for these devices, economics do not always incentivize the most secure choices.

1.2. Abuse examples

Two prominent abuse examples from the last two years demonstrate the threat posed by adversaries leveraging home Internet CPE: (1) distributed denial of service (DDoS) reflection and amplification and (2) man-in-the-middle (a.k.a. middleperson) attacks that redirect traffic from the local network by leveraging changes to the CPE to perform the middleperson attack.

A denial of service (DoS) attack is "the prevention of authorized access to resources or the delaying of time-critical operations."¹ A DDoS attack is executed from a distributed set of resources, making the attack harder to block.

¹ Kissel, Richard (ed.). Glossary of Key Information Security Terms. U.S. Department of Commerce – National Institute of Standards and Technology. NISTIR 7298, rev. 2. May 2013. <http://nvlpubs.nist.gov/nistpubs/ir/2013/NIST.IR.7298r2.pdf>

If the set of attackers is large and distributed such that it looks like regular traffic from customary users (like home users' routers, for example), the attack is almost impossible to block. However, creating this condition depends on some method by which the adversary can cause all these CPE devices to generate traffic to the intended victim.

Unfortunately, this traffic generation is often trivial. Due to the way DNS queries are addressed using UDP in the TCP/IP protocol stack, the source of any query is not authenticated. The adversary can forge the address of an intended victim in queries. Unless configured properly to only respond to users in the home or office, well-meaning DNS resolvers will think that the victim asked the question and then respond...and respond...and respond. A lot of home-Internet CPE is mistakenly configured to answer DNS queries from anyone on the Internet, which means the adversary can issue queries to hundreds or thousands of CPE devices, and they will dutifully all respond to the victim simultaneously, often causing a DoS condition. The insidious thing about this tactic is that it is reflected off the CPE, making attribution to the adversary nigh impossible while simultaneously causing the hard-to-filter traffic appear to come from regular users. Further, DNS queries are much smaller than DNS responses, so the adversary can increase attack volume by 20 times, using the home user's computer resources to amplify the attack. Amplification factors of up to 50 have been observed where the crypto-authentic DNS Security Extensions (DNSSEC) protocol was used. US-CERT issued an alert in March 2013 about this abuse of the DNS.² The DNS Operations Analysis and Research Center (DNS-OARC) summarizes the problem of why DNS is such an effective DDoS vector nicely:³

- *DNS generally uses the connectionless User Datagram Protocol (UDP) as its transport.*
- *Many autonomous systems allow source-spoofed packets to enter their networks.*
- *There is no shortage of open resolvers on the Internet.*

The other prevalent form of abuse that has affected CPE is middleperson attacks taking advantage of the router to target all users behind it in the home or business. In 2011, the FBI and several partners disrupted a botnet that was doing this on a large scale.⁴ This allowed the criminals to attack users whose computers they had not infected but who shared a home router with one that had been infected. Apparently the most lucrative use of this botnet was to redirect Internet-based advertising; they did this by redirecting all DNS queries from the victims to the criminals.⁵ A key method of redirecting the traffic was to abuse weak CPE and change the DNS settings in those devices. Once this was done, the adversaries controlled almost everywhere that their victims went on the

² United States Computer Emergency Readiness Team. Alert (TA13-088A) DNS Amplification Attacks. March 29, 2013; revised July 22 2013. <https://www.us-cert.gov/ncas/alerts/TA13-088A>

³ Domain Name System Operations Analysis and Research Center. "Mitigating DNS Denial of Service Attacks." <https://www.dns-oarc.net/wiki/mitigating-dns-denial-of-service-attacks>. Accessed Jun 16 2014.

⁴ United States Department of Justice, Federal Bureau of Investigation. "Operation Ghost Click: International Cyber Ring That Infected Millions of Computers Dismantled." Nov 9, 2011. http://www.fbi.gov/news/stories/2011/november/malware_110911.

⁵ Domain Changer Working Group. DCWG home page. <http://www.dcwg.org/>. Accessed Jun 18, 2014.

Internet. This could have been used for much more insidious purposes than advertising substitution.

2. Scope of the problem

There are several dimensions of the problem to define. First we define simply the number of devices connected to the Internet that fall under the category of CPE. Next we provide a general feel for how soft of a target these CPE devices are. The targets are soft targets due to the general attack surface, known vulnerabilities, and the rate at which these holes can be patched. In order to ground the problem, we select a particular example of CPE abuse that is ongoing—open DNS resolver reflection—and present measurements. First we measure the number of open resolvers and quantify approximately how many are CPE versus other kinds of Internet-connected equipment. Then we provide known measurements of the problems that DNS-based DDoS attacks have caused using open resolvers as reflectors.

Descriptions of hacks and control of CPE are more difficult to provide, as such operations are more clandestine. DNS changer provides a case study where the trojan allowed criminals to steal a net-total estimate of \$14 million by manipulating the DNS in part via changes to CPE router settings.⁶ This allowed the criminals to attack users whose computers they had not infected but who shared a home router with one that had been infected. Yet the underlying CPE vulnerabilities remain.

2.1. Statistics on CPE devices sold

Approximately 86.1 million US households are connected to broadband Internet, as of July 2013.⁷ In the United States, about 84% of those connections have a local home network, such as WiFi.⁸ Thus there are at least 72.5 million CPE routers and wireless routers just in the United States. Across the European Union, all 200 million households have broadband access.⁹ There are hundreds of millions of devices that can connect to the Internet through these local wireless connection points, with 324 million WiFi-enabled consumer electronics devices shipped during 2013 in the United States alone.¹⁰

Thus if there were a systematic vulnerability with CPE devices, in the United States or internationally, the impact on the public Internet would be pervasive just due to the sheer number of CPE devices. Their critical placement

⁶ United States Department of Justice, Federal Bureau of Investigation. “Operation Ghost Click: International Cyber Ring That Infected Millions of Computers Dismantled.” Nov 9, 2011. http://www.fbi.gov/news/stories/2011/november/malware_110911

⁷ IHS, Inc. “Broadband Internet Penetration Deepens in US; Cable is King.” Dec 9, 2013. <https://technology.ihs.com/468148/broadband-internet-penetration-deepens-in-us-cable-is-king>. Accessed 6/16/2014.

⁸ Broadband TV News. “84% of US broadband households have home network.” May 13, 2013. <http://www.broadbandtvnews.com/2013/05/13/84-of-us-broadband-households-have-home-network/>. Accessed 6/16/2014.

⁹ European Commission. Digital Agenda Scoreboard 2014 – Broadband Markets. May 28, 2014.

http://ec.europa.eu/information_society/newsroom/cf/dae/document.cfm?doc_id=5810

¹⁰ Watkins, David. Embedded WLAN (Wi-Fi) CE Devices: USA Market Forecast. Strategy Analytics. Feb 25, 2014. <http://www.strategyanalytics.com/default.aspx?mod=reportabstractviewer&a0=9405>

as the mediator between consumer devices and Internet access only amplifies the criticality of these devices. Unfortunately, despite the importance of the devices, the state of the practice for CPE security is well behind that for other computer equipment.

2.2. Vulnerability database information on CPE devices

A large number of CPE devices contain vulnerabilities, varying from simple cross-site scripting bugs to open ports on the Internet-facing-interface (i.e. the WAN interface). In particular, 27 remotely exploitable vulnerabilities were identified in CPE devices from June 2013 to June 2014.¹¹ We have noticed an uptick in vulnerability reports relating to these devices, including vulnerabilities related to suspected manufacturer or ISP-provided backdoors.¹²

The major challenge with these vulnerabilities is that the end user cannot patch them easily, unlike a traditional operating system or application patch. Routers require users to check for upgrades and then download and update the associated firmware manually. Another issue is that these devices are replaced infrequently, unlike a laptop or phone, so any vulnerabilities not patched by the user are left remnant on the device until it is replaced.

2.3. Measurement of open DNS resolvers

Open DNS resolvers are DNS services that will answer anyone's question. This openness is problematic because the source of the question can be easily forged, and the open resolvers can be used to create a DDoS attack via DNS reflection. There are quite a few open resolvers on the Internet. Although the number has decreased recently, as Figure 1 shows, it is still dangerously high at around 22.5 million.

The OpenResolverProject scans the Internet for open resolvers and makes the scan results available to security professionals. We now have 13 months of weekly scans. The general intuition is that these open resolvers are largely CPE, but that hasn't been quantified. If you need to check your own network for open resolvers, you can do so at the Measurement Factory website also.¹³

¹¹ NIST. National Vulnerability Database. www.nvd.nist.gov. Accessed 6/25/2014. Report numbers: CVE-2014-0356, CVE-2014-0354, CVE-2014-0353, CVE-2014-1982, CVE-2014-2925, CVE-2014-3792, CVE-2013-4772, CVE-2014-2719, CVE-2013-5948, CVE-2014-0337, CVE-2014-1599, CVE-2013-3365, CVE-2013-3098, CVE-2014-0329, CVE-2013-3090, CVE-2013-3087, CVE-2013-3084, CVE-2013-6343, CVE-2014-0659, CVE-2013-7282, CVE-2013-7043, CVE-2013-6918, CVE-2013-3095, CVE-2013-2271, CVE-2013-5703, CVE-2013-6027, CVE-2013-6026

¹² Gallagher, Sean. "Backdoor in wireless DSL routers lets attacker reset router, get admin." Jan 2, 2014. <http://arstechnica.com/security/2014/01/backdoor-in-wireless-dsl-routers-lets-attacker-reset-router-get-admin/>

¹³ The Measurement Factory. Open Resolver Test. <http://dns.measurement-factory.com/cgi-bin/openresolvercheck.pl>. Accessed Jun 16, 2014.




Figure 1: Number of open resolvers as reported by the OpenResolverProject.

We can estimate the type of devices that are abused by correlating the open resolver data with IP Intelligence data from Neustar (formerly Quova).¹⁴ That data identifies connection type and speed, assigning 1 of 13 different possible types to an IP address. Cable, dial-up, and DSL connections tend to be indicative of consumer connections, whereas optic, leased line, and other high-speed connections are indicative of businesses. Mobile wireless connections are consumer devices, and as smart phones become more complex, they will need to follow the same recommendations as for CPE. However we are generally not considering cellular devices as the same kind of problem.

The open resolvers show a markedly different and statistically significant pattern of connection types than on the Internet at-large. The first step is to establish a baseline. IP address usage does change over time, but as demonstrated by in Figure 2, it is stable. A large portion of the address space is unassigned and unused, and there are many addresses that the IP intelligence does not label for speed. But, the next largest contingent is the roughly 18% of IP space comprising leased lines for medium- and large-sized companies.

¹⁴ Neustar. GeoPoint Data Glossary. Accessed May 28, 2014.
<https://ipintelligence.neustar.biz/portal/#documentation>




Figure 2: Connection type and speed of every IP addresses across the Internet on three dates.

Compare this baseline with the trends in open resolvers between April 2013 and May 2014, as displayed in Figure 3. DSL makes up between 45-50% of the open resolver IP addresses at every measurement, compared to 11% in the baseline frequency of the Internet. This is a striking difference; the results are statistically significant and not due to random chance. See Appendix B for statistical details. Open resolvers are much more likely to be on home and small-business customer lines, and much less likely to be on leased, large-business lines. Further, the only way a DNS resolver should be open on a DSL connection is due to an error in the CPE configuration.

Thus not only are there quite a few home routers and pieces of CPE in the world, as reported in Section 2.1, but devices on customer premises are disproportionately available for abuse in creating reflected and amplified DDoS attacks that can threaten the whole Internet community. If each DSL- and cable-connected open resolver could sustain even a modest 1 Mbps output, the 11 million devices would produce roughly 1.1 terabits per second. This amount of traffic would not necessarily disrupt the core of the Internet; however 1.1 Tbps could disrupt a large ISP.



Figure 3: Number of open resolvers on the public Internet and their connection type.

2.4. DDoS attacks using open resolvers

The high watermark for the largest DDoS traffic volume openly reported as of 2013 was the result of reflection off open resolvers. The attack on Spamhaus in March 2013 measured a maximum of about 300 Gbps.¹⁵ Since then, CloudFlare has stated there have been larger reflection attacks using network time protocol's (NTP) MONLIST operator. However these attacks are reported at only a coarse detail level, and the actual largest single attack may not have even been publicly reported.¹⁶ Open resolvers have also been used in sizable

¹⁵ Prince, Matthew. "The DDoS That Almost Broke the Internet." CloudFlare. Mar 27, 2013. <http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet>. Accessed Jun 19, 2014.

¹⁶ Greenberg, Adam. CloudFlare fights off massive NTP reflection DDoS attack. SC Magazine. Feb 11, 2014. <http://www.scmagazine.com/cloudflare-fights-off-massive-ntp-reflection-ddos-attack/article/333585/>

DDoS attacks against real-time financial exchange platforms.¹⁷ In addition, open resolvers are used as an attack amplification method regularly for smaller attacks that largely go unreported, and it appears that attackers are now crafting DNS zones specifically to make maximum use of this threat's amplification abilities.¹⁸ Although a Neustar DDoS survey does not ask about attack technique, it is telling that 60% of companies surveyed suffered at least one DDoS attack of some kind, and most of those were attacked more than once.¹⁹ Neustar did see a rise in large (100+ Gbps) DDoS attacks, which it attributes largely to the rise in DNS and NTP reflection attacks.

3. Recommendations

We suggest recommendations in three areas: continuous upgrade path for all devices, source address validation, and clearer responsibility for failings while aligning incentives.

3.1. Device manufacturers need a path for continuous upgrades

Software support for CPE is far from cradle-to-grave. Indeed we can better characterize it as fire-and-forget, insofar as CPE manufacturers often have no plan for ever revising the software the device shipped with and no method of distributing or installing such patches. Error theory tells us that all software has bugs, so the absence of a patch distribution mechanism for devices expected to be sold in the tens of millions of units is at best shortsighted. A device that works well when it works but is a danger to the Internet commons when misused must be patchable in the field. Obviously, care must be taken to ensure that the patch distribution and installation data path cannot be hijacked to distribute malware to CPE. Traditional code-signing techniques should work in this application.

3.2. Implement source address validation

Most of the Internet's edge is not secure in the sense that forged IP source addresses are allowed to propagate toward the Internet's core, unnoticed and unchallenged. A small number of corner cases such as multi-homing and multi-pathing require a gateway to accept source addresses from an edge network. Those cases should be treated as exceptions and require non-default configurations of both the customer-premise and ISP equipment. Such treatment would add the small number of alternate network addresses as permitted IP sources that might actually be necessary in multi-homed or multi-path configurations. Prevention of IP source address forgery is called *source address*

¹⁷ Prolexic. "Prolexic Stops Largest-Ever DNS Reflection DDoS Attack." May 30, 2013. <http://www.prolexic.com/news-events-pr-prolexic-stops-largest-ever-dns-reflection-ddos-attack-167-gbps.html>. Accessed Jun 19, 2014.

¹⁸ Weber, Ralph. Better than Best Practices For DNS Amplification Attacks. NANOG 59th meeting, Phoenix, AZ, October 2013. https://www.nanog.org/sites/default/files/mon_general_weber_defeat_23.pdf

¹⁹ Neustar. 2014 -- The Danger Deepens: Neustar Annual DDoS Attacks and Impact Report. 2014. <http://www.neustar.biz/resources/whitepapers/ddos-protection/2014-annual-ddos-attacks-and-impact-report.pdf>

validation (S.A.V.), which must become the permanent worldwide default configuration for Internet traffic-carrying devices

Importantly, there are two entirely different Internet edges where S.A.V. is relevant, and only one of those edges is the CPE device pool that is the main topic of this report. For CPE devices, the customer's visible network most often uses non-routable addresses as described in RFC 1918, and one of the CPE's responsibilities is to translate these addresses into routable addresses when packets propagate from the edge toward the core. Some CPE devices offer an unfortunate loophole in which IP packets already having routable IP source addresses are propagated unchanged. Thus an attack launched by a malware-infected personal computer inside the customer's edge network can forge the routable IP address of its intended victim, knowing that those addresses will not be subject to the same Network Address Translation (NAT) as the customer's own benign traffic.

The other relevant network edge where S.A.V. is almost universally not practiced today is "the cloud," including providers of either physical or virtual server hosting. Hosting providers operate on very thin margins, and the added cost of configuring and auditing S.A.V. in their customer-facing gateways would apparently threaten their profitability. Since the only beneficiary of spending on S.A.V. would be the Internet commons—in other words, not their own customers—it is extremely difficult to justify this expense. Nevertheless for the continued health and growth of the Internet ecosystem, these costs must be accepted by these operators.

3.3. Responsibility

One certain path to slow poisoning of an ecosystem is for every participant to regard long-term global challenges as *that's not my problem*. Indeed the response to carbon-induced climate change, to de-oxygenation caused by excess fertilizer use and subsequent runoff, to the evolution of antibiotic-resistant strains of bacteria, and to a hundred other long-term global challenges facing humanity is a near-universal claim by every actor in the system to the effect that *nothing I can do will make a measurable difference*. Game theory backs up this assertion: It is statistically better for an individual to act from irrational short-term selfishness than to cooperate. However, that analysis presumes incomplete knowledge, which can only be the case here if the actors widely practice self-deception and denial. The long-term rational self-interest of every actor on the Internet is to preserve the health and growth of the Internet, as a benefit to every actor in the long term.

The misalignment of incentives among end users, manufacturers, and operators is essentially due to the fact that the problems are externalities that can be put off in the short term. The usual response to controlling externalities that threaten a community is regulation of how they can be handled. Reining in the observed widespread Internet abuse would likely require equally widespread regulation among the world's democracies, in part to ensure a healthy and growing Internet, and in part to ensure that compliance costs are imposed universally. Such regulation would prevent a hefty disincentive in the form of cost structures not shared by all competitors for those acting in the best interest of the Internet commons.

4. Conclusion

Misconfigured, poorly thought-out CPE choices remain a threat to the continued health of the public Internet. As more consumer devices connect to the Internet and fewer of them have human interfaces, the Internet community must understand that the decisions of device manufacturers and network engineers matter to everyone, not just the direct users of those devices or services.

The scope of this problem is large, and the solutions are difficult. In many ways, the problem is analogous to the public health issues faced as cities grow large and crowded. The dangerous habits of relatively small populations can threaten the larger population in complex ways that must be investigated, and the community must respond to manage these dangerous habits. Given the number of known bugs and poor update cycles on CPE as described in Section 2.2, this poor device “health” in the deployed CPE population should come as no surprise.

The situation is pressing but not hopeless; there are well-defined solutions the community can work toward. Software upgrades and source address validation are not new practices. The important next step is the uniform adoption of such minimum standards led by the will of the community and responsible manufacturers to force all relevant parties to take responsibility for these standards.

Bringing about this change would require a concerted effort from many parties and take several years. And the community must not forget these lessons as new classes of devices are inevitably invented and connected to the network. However it is critical that the work be continued through to the end, or the Internet will remain a dangerous place where connectivity can be interrupted essentially at the whim of an adversary bouncing attacks off unwittingly misconfigured CPE.

Appendix A – Neustar Connection types and descriptions

Quoted from the Neustar IP Intelligence documentation at
<https://ipintelligence.neustar.biz/portal/#documentation>

cable	<i>Cable Modem broadband circuits, offered by cable TV companies. Speeds range from 128 Kbps to 100 Mbps, and vary with the load placed on a given cable modem switch.</i>
consumer satellite	<i>High-speed or broadband links between a consumer and a geosynchronous or low-earth orbiting satellite. By default, IP addresses with a consumer satellite Connection Type are assigned a satellite IP Routing Type as well. See the satellite IP Routing Type for more information.</i>
dial-up	<i>Consumer dial-up modem technology, which operates at 56 Kbps. Providers include Earthlink, AOL, and Netzero.</i>
dsl	<i>Digital Subscriber Line broadband circuits, which include aDSL, iDSL, sDSL, etc. DSL ranges in speed from 256 Kbps (kilobits per second) to 20 Mbps (megabits per second).</i>
fixed wireless	<i>Fixed wireless connections, where the location of the receiver is fixed. This category includes WDSL providers such as Sprint Broadband Direct, as well as emerging WiMax providers.</i>
framerelay	<i>Frame relay circuits, which can range from low- to high-speed and are used as a backup or alternative to T-1. Most often, they are high-speed links, so IP Intelligence classifies them as such.</i>
isdn	<i>Integrated Services Digital Network high-speed copper-wire technology, which provides 128 Kbps speed, with ISDN modems and switches offering 1 Mbps and greater speeds. Offered by some major telephony companies.</i>
mobile wireless	<i>Cellular network providers such as AT&T, Sprint, and Verizon Wireless who employ CDMA, EDGE, EV-DO, GPRS, 3G, and 4G technologies. Speeds vary from 19.2 Kbps to 3 Mbps.</i>
ocx	<i>Fiber optic connections (including OC-3, OC-48, OC-192, etc.), which are used primarily by large backbone carriers.</i>
tx	<i>Leased line, that is, T1, T2, T3, or T4, circuits used by many small- and medium-sized companies.</i>
unknown high	<i>Indicates that IP Intelligence was unable to obtain the connection type. However, the estimated connection speed is high.</i>
unknown low	<i>Indicates that IP Intelligence was unable to obtain the connection type. However, the estimated connection speed is low.</i>
unknown medium	<i>Indicates that IP Intelligence was unable to obtain the connection type. However, the estimated connection speed is medium.</i>

Appendix B – Statistical tests

We performed a χ^2 independence test over the categorizations of connections types (DSL, TX, etc.) for the baseline set (whole Internet) and open resolver IP sets based on the number of IP addresses in each category using the data for May 2014. May 11 is the date of the open resolver scan, and May 9 is the date of connection-type markings. This yielded a P value of 0, indicating that the distribution over the connection type categories was dependent on whether the IP was an open resolver. Then we performed a pairwise conditional dependence test for each category of connection type between the two measurements with a 0.99 confidence interval for the range of expected values over the open resolver measurements. Pairwise tests for all connection types are statistically significant at this value: The 0.99 confidence interval is less than the difference between the ratios.

Table 1: Confidence interval and difference (in percentage points) between measurements for open resolvers and the Internet at-large on May 11 and May 9 2014 measurements, respectively. The percentage of IPs labeled as open resolvers for each connection type is p_1 ; the percentage labeled as the connection type on the Internet at-large is p_2 . In all cases, the confidence interval is orders of magnitude smaller than the difference.

Connection type	.99 Confidence interval (\pm)	$\hat{p}_1 - \hat{p}_2$
Speed unknown	0.00000566%	5.927%
tx-high	0.00000232%	-14.452%
Unassigned	0.00000033%	-19.265%
ocx-high	0.00000031%	-0.061%
dial-up-low	0.00000080%	-0.944%
dsl-medium	0.00000573%	35.177%
isdn-medium	0.00000012%	-0.038%
cable-medium	0.00000222%	-2.603%
framerelay-high	0.00000015%	-0.043%
unknown high-high	0.00000002%	-0.001%
mobile wireless-low	0.00000201%	-3.617%
fixed wireless-medium	0.00000075%	0.033%
unknown medium-medium	0.00000004%	-0.002%
consumer satellite-medium	0.00000018%	-0.112%